

知 CAS management web page access abnormal performance, data refresh is slow

Cloud Computing 厉梦如 2020-09-27 Published

Network Topology

Version: CAS E0530H12

Problem Description

Access to CVM page is abnormal, the performance data of virtual machine interface is very slow to refresh, the storage pool page has been circling, and the overall operation of the platform is stuck.

Process Analysis

1. First of all, excluding the network and environmental factors: check whether the client is two-layer access or cross-three-tier, whether there is a problem with the client computer, whether the phenomenon of trying to access with other PCs have the same problem;
2. Check the number of Tomcat connections: ssh to CVM node ,then use command `netstat - NAT | grep 8080 | WC - L`. If it turns out the number of connections is too large, then may cause the page load be stuck. In normal environment, the number of connections is about 200, which is normal and has no fixed value;
3. Check whether the storage is normal: Perform iostat to check the storage delay and busy degree. It is found that the % util of dm-0 storage volume reaches 100%, the disk is busy, there are many I/O requests, and it runs at full load.

Device:	rrqm/s	wrqm/s	r/s	w/s	rkB/s	wkB/s	avgqz-sz	avgqu-sz	await	r_await	w_await	svctm	%util
dm-0	0.00	0.00	0.00	0.50	0.00	1.25	5.00	6.00	4.00	0.00	4.00	2000.00	100.00
Device:	rrqm/s	wrqm/s	r/s	w/s	rkB/s	wkB/s	avgqz-sz	avgqu-sz	await	r_await	w_await	svctm	%util
dm-0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	6.00	0.00	0.00	0.00	0.00	100.00
Device:	rrqm/s	wrqm/s	r/s	w/s	rkB/s	wkB/s	avgqz-sz	avgqu-sz	await	r_await	w_await	svctm	%util
dm-0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	6.00	0.00	0.00	0.00	0.00	100.00
Device:	rrqm/s	wrqm/s	r/s	w/s	rkB/s	wkB/s	avgqz-sz	avgqu-sz	await	r_await	w_await	svctm	%util
dm-0	0.00	0.00	2.50	3.50	18.25	14.50	10.82	5.41	14014.33	17966.40	11191.43	166.67	100.00

4. Confirm whether there is cluster lock :ssh to cvk shell line and execute `OCFS2_kernel_stat_colle ct.sh` command to check whether there is a hung up process.As shown in the figure below, it is found that there is a hung up process and dm-0 is busy printing. A similar log file named `cvk01_ocfs2_kernel_state_202005231721.log` is generated in the / var / log directory. The log file contains the following logs. It is preliminarily judged that IO is blocked

```
==== find hung up process =====
4058 D fault_checker Mon May 11 11:05:55 2020 msleep
35447 D debugfs.ocfs2 Sat May 23 17:44:45 2020 blkdev_direct_IO

==== get stack of 4058 =====

[<ffffffff832fa6a9>] msleep+0x29/0x30
[<ffffffffc09a23d9>] o2quo_check_fault_thread+0x49/0x120 [ocfs2_nodemanager]
[<ffffffff832a46ea>] kthread+0x11a/0x130
[<ffffffff83c00205>] ret_from_fork+0x35/0x40
[<fffffffffffffff>] 0xffffffffffff

==== get stack of 35447 =====
debugfs.ocfs2 -R hb /dev/dm-0
[<ffffffff832b16b2>] io_schedule+0x12/0x40
[<ffffffff83493acf>] __blkdev_direct_IO_simple+0x1ff/0x360
[<ffffffff83493fd2>] blkdev_direct_IO+0x3a2/0x3f0
[<ffffffff833b6af7>] generic_file_read_iter+0xb7/0xb90
[<ffffffff83455edf>] __vfs_read+0xff/0x170
[<ffffffff83455fd9>] vfs_read+0x89/0x130
[<ffffffff83456720>] Sys_pread64+0x90/0xb0
[<ffffffff832039ce>] do_syscall_64+0x6e/0x120
[<ffffffff83c00081>] entry_SYSCALL_64_after_hwframe+0x3d/0xa2
[<fffffffffffffff>] 0xffffffffffff
```

5. Check the syslog log log:there are a large number of links concussion, multi-path connection failure print, so we can judge the storage failure

```

23 18:03:45 cvk02 kernel: [37766.632488] sd 2:0:0:0: [sdc] tag#0 Add. Sense: Logical unit not accessible, target port in standby state
23 18:03:45 cvk02 kernel: [37766.632492] sd 2:0:0:0: [sdc] tag#0 CDB: Read(16) 88 00 00 00 00 00 00 28 00 00 00 00 11 00 00
23 18:03:45 cvk02 kernel: [37766.632494] print_req_error: I/O error, dev sdc, sector 10240
23 18:03:45 cvk02 kernel: [37766.632505] device-mapper: multipath: Failing path 8:32.
23 18:03:45 cvk02 multipathd: reinsstate_path:1372: 8:16: reinstated
23 18:03:45 cvk02 multipathd: update_queue_mode_add_path:630: 360002ac0000000000000000300024b13: remaining active paths: 4
23 18:03:45 cvk02 multipathd: update_multipath:571: sdc: mark as failed
23 18:03:45 cvk02 multipathd: update_queue_mode_del_path:618: 360002ac0000000000000000300024b13: remaining active paths: 3
23 18:03:45 cvk02 multipathd: update_multipath:571: sdd: mark as failed
23 18:03:45 cvk02 multipathd: update_queue_mode_del_path:618: 360002ac0000000000000000300024b13: remaining active paths: 2
23 18:03:45 cvk02 kernel: [37767.438646] device-mapper: multipath: Reinstating path 8:16.
23 18:03:45 cvk02 multipathd: reinsstate_path:1372: 8:32: reinstated
23 18:03:45 cvk02 multipathd: update_queue_mode_add_path:630: 360002ac0000000000000000300024b13: remaining active paths: 3
23 18:03:45 cvk02 kernel: [37771.439854] device-mapper: multipath: Reinstating path 8:32.
23 18:03:45 cvk02 multipathd: reinsstate_path:1372: 8:48: reinstated
23 18:03:45 cvk02 multipathd: update_queue_mode_add_path:630: 360002ac0000000000000000300024b13: remaining active paths: 4
23 18:03:45 cvk02 kernel: [37771.440509] device-mapper: multipath: Reinstating path 8:48.
23 18:03:45 cvk02 multipathd: update_multipath:571: sdc: mark as failed
23 18:03:45 cvk02 multipathd: update_queue_mode_del_path:618: 360002ac0000000000000000300024b13: remaining active paths: 3
23 18:03:45 cvk02 multipathd: update_multipath:571: sdd: mark as failed
23 18:03:45 cvk02 multipathd: update_queue_mode_del_path:618: 360002ac0000000000000000300024b13: remaining active paths: 2
23 18:03:45 cvk02 kernel: [37772.680480] sd 1:0:1:0: [sdd] tag#0 FAILED Result: hostbyte=DID_OK driverbyte=DRIVER_SENSE
23 18:03:45 cvk02 kernel: [37772.680486] sd 1:0:1:0: [sdd] tag#0 Sense Key : Not Ready [current]
23 18:03:45 cvk02 kernel: [37772.680491] sd 1:0:1:0: [sdd] tag#0 Add. Sense: Logical unit not accessible, target port in standby state
23 18:03:45 cvk02 kernel: [37772.680497] sd 1:0:1:0: [sdd] tag#0 CDB: Read(16) 88 00 00 00 00 00 00 28 00 00 00 00 11 00 00
23 18:03:45 cvk02 kernel: [37772.680501] print_req_error: I/O error, dev sdd, sector 10240
23 18:03:45 cvk02 kernel: [37772.680523] device-mapper: multipath: Failing path 8:48.
23 18:03:45 cvk02 kernel: [37772.680634] sd 2:0:0:0: [sdc] tag#0 FAILED Result: hostbyte=DID_OK driverbyte=DRIVER_SENSE
23 18:03:45 cvk02 kernel: [37772.680637] sd 2:0:0:0: [sdc] tag#0 Sense Key : Not Ready [current]
23 18:03:45 cvk02 kernel: [37772.680641] sd 2:0:0:0: [sdc] tag#0 Add. Sense: Logical unit not accessible, target port in standby state
23 18:03:45 cvk02 kernel: [37772.680646] sd 2:0:0:0: [sdc] tag#0 CDB: Read(16) 88 00 00 00 00 00 00 28 00 00 00 00 11 00 00

```

6. Multipath - ll view multipath, with failed print

```

root@cvk02:~# multipath -ll
360002ac0000000000000000300024b13 dm-0 3PARdata,VV
size=10T features='l queue if no path' hwhandler='0' wp=rw
+- policy='round-robin 0' prio=25 status=active
|- 1:0:0:0 sdb 8:16 failed ready running
|- 2:0:0:0 sdc 8:32 failed ghost running
|- 1:0:1:0 sdd 8:48 failed ghost running
+- 2:0:1:0 sde 8:64 active ready running:

```

Solution

Check the storage and storage link conditions as soon as possible according to the corresponding storage manufacturer, such as optical power, packet loss, etc., and replace the HBA card or optical fiber, optical port, etc. if necessary.

附件下载: CAS management web page access abnormal performance